

TITLE OF THE INVENTION:

[0001] SWITCH HAVING EXTERNAL ADDRESS RESOLUTION INTERFACE

REFERENCE TO RELATED APPLICATIONS:

[0002] This application claims priority to United States Provisional Patent Application Serial No. 60/213,478 filed on June 23, 2000 and United States Provisional Patent Application Serial No. 60/213,482 filed on June 23, 2000. The contents of the provisional applications identified above are hereby incorporated by reference.

BACKGROUND OF THE INVENTION:**Field of the Invention:**

[0003] The invention relates to a method and apparatus for high performance switching in local area communications networks such as token ring, ATM, ethernet, fast ethernet, and gigabit ethernet environments, generally known as LANs. In particular, the invention relates to a new switching architecture in an integrated, modular, single chip solution, which can be implemented on a semiconductor substrate such as a silicon chip.

Description of the Related Art:

[0004] As computer performance has increased in recent years, the demands on computer networks has significantly increased; faster computer processors and higher memory capabilities need networks with high bandwidth capabilities to enable high speed transfer of significant amounts of data. The well-known ethernet technology, which is based upon numerous IEEE ethernet standards, is one example of computer

networking technology which has been able to be modified and improved to remain a viable computing technology. A more complete discussion of prior art networking systems can be found, for example, in SWITCHED AND FAST ETHERNET, by Breyer and Riley (Ziff-Davis, 1996), and numerous IEEE publications relating to IEEE 802 standards. Based upon the Open Systems Interconnect (OSI) 7-layer reference model, network capabilities have grown through the development of repeaters, bridges, routers, and, more recently, "switches", which operate with various types of communication media. Thickwire, thinwire, twisted pair, and optical fiber are examples of media which has been used for computer networks. Switches, as they relate to computer networking and to ethernet, are hardware-based devices which control the flow of data packets or cells based upon destination address information which is available in each packet. A properly designed and implemented switch should be capable of receiving a packet and switching the packet to an appropriate output port at what is referred to as wirespeed or linespeed, which is the maximum speed capability of the particular network. Basic ethernet wirespeed is up to 10 megabits per second, and Fast Ethernet is up to 100 megabits per second. The newest ethernet is referred to as gigabit ethernet, and is capable of transmitting data over a network at a rate of up to 1,000 megabits per second. As speed has increased, design constraints and design requirements have become more and more complex with respect to following appropriate design and protocol rules and providing a low cost, commercially viable solution. For example, high speed switching requires high speed memory to provide appropriate buffering of packet

data; conventional Dynamic Random Access Memory (DRAM) is relatively slow, and requires hardware-driven refresh. The speed of DRAMs, therefore, as buffer memory in network switching, results in valuable time being lost, and it becomes almost impossible to operate the switch or the network at linespeed. Furthermore, external CPU involvement should be avoided, since CPU involvement also makes it almost impossible to operate the switch at linespeed. Additionally, as network switches have become more and more complicated with respect to requiring rules tables and memory control, a complex multi-chip solution is necessary which requires logic circuitry, sometimes referred to as glue logic circuitry, to enable the various chips to communicate with each other. Additionally, cost/benefit tradeoffs are necessary with respect to expensive but fast SRAMs versus inexpensive but slow DRAMs. Additionally, DRAMs, by virtue of their dynamic nature, require refreshing of the memory contents in order to prevent losses thereof. SRAMs do not suffer from the refresh requirement, and have reduced operational overhead which compared to DRAMs such as elimination of page misses, etc. Although DRAMs have adequate speed when accessing locations on the same page, speed is reduced when other pages must be accessed.

[0005] Referring to the OSI 7-layer reference model discussed previously, and illustrated in Figure 7, the higher layers typically have more information. Various types of products are available for performing switching-related functions at various levels of the OSI model. Hubs or repeaters operate at layer one, and essentially copy and "broadcast" incoming data to a plurality of spokes of the hub. Layer two switching-

related devices are typically referred to as multiport bridges, and are capable of bridging two separate networks. Bridges can build a table of forwarding rules based upon which MAC (media access controller) addresses exist on which ports of the bridge, and pass packets which are destined for an address which is located on an opposite side of the bridge. Bridges typically utilize what is known as the "spanning tree" algorithm to eliminate potential data loops; a data loop is a situation wherein a packet endlessly loops in a network looking for a particular address. The spanning tree algorithm defines a protocol for preventing data loops. Layer three switches, sometimes referred to as routers, can forward packets based upon the destination network address. Layer three switches are capable of learning addresses and maintaining tables thereof which correspond to port mappings. Processing speed for layer three switches can be improved by utilizing specialized high performance hardware, and off loading the host CPU so that instruction decisions do not delay packet forwarding.

SUMMARY OF THE INVENTION:

[0006] The invention is directed to a network switch that has a plurality of input ports that receive data packets. An external address resolution interface is connected to the plurality of input ports. The external address resolution interface externally transmits the data packets for processing, and receives the data packets after processing. A memory management unit is connected to the external interface and a plurality of output ports are connected to the memory management unit.

[0007] The invention is also directed to a method of processing a data packet in a network switch. The method has the steps of receiving a data packet in an input port and transmitting the data packet from the input port over an interface to an external module for address resolution. The packet is then processed in the external module and transmitted from the external module to the interface. The packet is then received in the interface from said external module and transmitted from the interface to a memory management unit. The packet is finally transmitted from the memory management unit to an output port.

[0008] Another embodiment of the invention is network switch having an input port receiving means for receiving a data packet in an input port; an input port transmitting means for transmitting said data packet from said input port over an interface to an external module for address resolution; a processing means for processing said packet in said external module; an external module transmitting means for transmitting said packet from said external module to said interface; an interface receiving means for receiving said data packet in said interface from said external module; an interface transmitting means for transmitting said data packet from said interface to a memory management unit; and a memory unit transmitting means for transmitting said data packet from said memory management unit to an output port.

BRIEF DESCRIPTION OF THE DRAWINGS:

[0009] The objects and features of the invention will be more readily understood with reference to the following description and the attached drawings, wherein:

Figure 1 is an illustration of a pGS switch of the present invention;

Figure 2 is an illustration of two pGS switches being connected via 10-GE Interfaces;

Figure 3 is an illustration of an fGS fabric;

Figure 4 is an illustration of eight pGS switches connected to one another through an fGS fabric; and

Figure 5 is a flow diagram of how a packet is transferred through a pGS of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS:

[0010] FIG. 1 is an illustration of a port gigabit switch (pGS) 100. pGS 100 has eight gigabit ports 102, a 10-GE (Gigabit Ethernet) Interface 104 and an Address Resolution Gigabit Switch (aGS) Interface 106.

[0011] pGS 100 is a Layer 2 device with on-chip Media Access Control (MAC) address tables and a Fast Filtering Processor (FFP). The device can operate in standalone mode or can be connected to other devices including other pGS devices or an external aGS to perform Layer 2, Layer 3 and FFP processing. pGS 100 is a single chip non-blocking switching solution for eight gigabit Ethernet ports and a 10-GE link. The FFP provides accelerated packet forwarding and enhanced packet flow through the use of functions such as classification, metering, shaping and re-marking.

[0012] Each gigabit port 102 of pGS 100 is associated with a corresponding MAC having an ingress for receiving data packets and an egress for transmitting data

[0013] In one embodiment of the invention each of the gigabit ports 102 can be connected to a CPS channel. Arbitration for the CPS channel occurs out of band. The CPS channel is monitored by each of the gigabit ports 102, and matching destination ports respond to appropriate transactions. C-channel arbitration is a demand priority round robin arbitration mechanism. If no requests are active, however, the default module, which can be selected during the configuration of pGS 100 can park on the channel and have complete access thereto. If all requests are active, the configuration of pGS 100 is such that a Pipeline Memory Management Unit (PMMU) is granted access every other cell cycle, and gigabit ports 102 share equal access to the C-channel on a round robin basis. The PMMU and gigabit ports 102 alternate access, and the PMMU gains access every other cycle.

[0014] In the protocol or P-channel a plurality of messages can be placed on the P-channel in order to properly direct flow of data flowing on the C-channel. Since the P-channel can be typically 32 bits wide, and a message typically requires 128 bits, four smaller 32 bit messages are put together in order to form a complete P-channel message. The following list identifies the fields and function and the various bit counts of the 128 bit message on the P-channel.

- [0015] **Opcode** - 2 bits long - Identifies the type of message present on the C channel 81;
- [0016] **IP Bit** - 1 bit long - This bit is set to indicate that the packet is an IP switched packet;
- [0017] **IPX Bit** - 1 bit long - This bit is set to indicate that the packet is an IPX switched packet;
- [0018] **Next Cell** - 2 bits long - A series of values to identify the valid bytes in the corresponding cell on the C channel 81;
- [0019] **SRC DEST Port** - 6 bits long - Defines the port number which sends the message or receives the message, with the interpretation of the source or destination depending upon Opcode;
- [0020] **Cos** - 3 bits long - Defines class of service for the current packet being processed;
- [0021] **J** - 1 bit long - Describes whether the current packet is a jumbo packet;
- [0022] **S** - 1 bit long - Indicates whether the current cell is the first cell of the packet;
- [0023] **E** - 1 bit long - Indicates whether the current cell is the last cell of the packet;
- [0024] **CRC** - 2 bits long - Indicates whether a Cyclical Redundancy Check (CRC) value should be appended to the packet and whether a CRC value should be regenerated;
- [0025] **P Bit** - 1 bit long - Determines whether MMU should Purge the entire packet;
- [0026] **Len** - 7 bytes - Identifies the valid number of bytes in current transfer;

[0027] - 2 bits - Defines an optimization for processing by the CPU 52; and

[0028] Bc/Mc Bitmap - 28 bits - Defines the broadcast or multicast bitmap.

Identifies egress ports to which the packet should be set, regarding multicast and broadcast messages.

[0029] Untag Bits/Source Port - 28/5 bits long - Depending upon Opcode, the packet is transferred from Port to MMU, and this field is interpreted as the untagged bit map. A different Opcode selection indicates that the packet is being transferred from MMU to egress port, and the last six bits of this field is interpreted as the Source Port field. The untagged bits identifies the egress ports which will strip the tag header, and the source port bits identifies the port number upon which the packet has entered the switch;

[0030] U Bit - 1 bit long - For a particular Opcode selection (0x01, this bit being set indicates that the packet should leave the port as Untagged; in this case, tag stripping is performed by the appropriate MAC;

[0031] CPU Opcode - 18 bits long - These bits are set if the packet is being sent to the CPU for any reason. Opcodes are defined based upon filter match, learn bits being set, routing bits, destination lookup failure (DLF), station movement, etc;

[0032] Time Stamp - 14 bits - The system puts a time stamp in this field when the packet arrives, with a granularity of 1 μ sec.

[0033] The opcode field of the P-channel message defines the type of message currently being sent. While the opcode is currently shown as having a width of 2 bits,

the opcode field can be widened as desired to account for new types of messages as may be defined in the future.

[0034] An early termination message is used to indicate that the current packet is to be terminated. During operation, as discussed in more detail below, the status bit (S) field in the message is set to indicate the desire to purge the current packet from memory. Also in response to the status bit all applicable egress ports would purge the current packet prior to transmission.

[0035] The Src Dest Port field of the P-channel message, as stated above, define the destination and source port addresses, respectively. Each field is 6 bits wide and therefore allows for the addressing of sixty-four ports.

[0036] The CRC field of the message is two bits wide and defines CRC actions. Bit 0 of the field provides an indication whether the associated egress port should append a CRC to the current packet. An egress port would append a CRC to the current packet when bit 0 of the CRC field is set to a logical one. Bit 1 of the CRC field provides an indication whether the associated egress port should regenerate a CRC for the current packet. An egress port would regenerate a CRC when bit 1 of the CRC field is set to a logical one. The CRC field is only valid for the last cell transmitted as defined by the E bit field of P-channel message set to a logical one.

[0037] As with the CRC field, the status bit field (st), the Len field, and the Cell Count field of the message are only valid for the last cell of a packet being transmitted as defined by the E bit field of the message.

[0038] Last, the time stamp field of the message has a resolution of 1 μ s and is valid only for the first cell of the packet defined by the S bit field of the message. A cell is defined as the first cell of a received packet when the S bit field of the message is set to a logical one value.

[0039] As is described in more detail below, the C channel and the P channel are synchronously tied together such that data on C channel is transmitted over the CPS channel while a corresponding P channel message is simultaneously transmitted.

[0040] The S channel is a 32-bit wide channel which provides a separate communication path within the pGS 100. The S channel is used for management by a Central Processing Unit (CPU), pGS 100 internal flow control, and pGS 100 inter-module messaging. The S channel is a sideband channel of the CPS channel, and is electrically and physically isolated from the C channel and the P channel. It is important to note that since the S channel is separate and distinct from the C channel and the P channel, operation of the S channel can continue without performance degradation related to the C channel and P channel operation. Conversely, since the C channel is not used for the transmission of system messages, but rather only data, there is no overhead associated with the C channel and, thus, the C channel is able to free-run as needed to handle incoming and outgoing packet information.

[0041] The S channel of the CPS channel provides a system wide communication path for transmitting system messages throughout pGS 100.

[0043] Opcode - 6 bits long - Identifies the type of message present on the S channel;

[0045] Src Port -6 bits long - Defines the port number of which the current S channel message originated;

[0047] C bit - 1 bit long - Logically defines whether the current S channel message is intended for the CPU 52.

[0049] DataLen - 7 bits long - Defines the total number of data bytes in the **Data** field;

[0051] Address - 32 bits long - Defines the memory address associated with the current command as defined in **opcode**;

[0052] Data - 0-127 bits long - Contains the data associated with the current opcode.

[0053] With the configuration of the CPS channel as explained above, the decoupling of the S channel from the C channel and the P channel is such that the bandwidth on the C channel can be preserved for cell transfer, and that overloading of the C channel does not affect communications on the sideband channel. An aGS interface 106 can be connected to a CPS channel.

[0054] The aGS Interface 106 can be a high speed serial interface or a parallel interface in order to transfer data as quickly as possible. aGS Interface 106 can be connected to a CPS channel at one end and on another end to an external aGS which performs Address Resolution Logic (ARL) functions on the external aGS. The bandwidth of the aGS interface can be such that it can transmit and receive packets from all ports on pGS 100. In addition when pGS 100 utilizes an external aGS using aGS Interface 102, the FFP provided on pGS 100 is not used and the FFP provided on the external aGS is utilized. The external aGS, in addition to performing ARL functions externally, can also be provided with a means for very low buffering to speed up processing. In some cases a FIFO is provided which will only result in a FIFO latency.

[0055] The FFP provided on pGS 100 in one embodiment of the invention has 16 filter masks (inclusive and exclusive) and a Rules Table having a depth of 128K. The FFP on the external aGS can have 32 filter masks and a Rules Table having a depth of 256K. From this example it is evident that one advantage of having an external aGS is that it allows for more flexibility and modularity in implementing ARL functions on an as

needed basis. If the on chip filter masks are not adequate for a given task, different filter masks can be used by implementing an external aGS using aGS Interface 106.

[0056] FIG. 2 illustrates a first pGS 210 having eight gigabit ports 212, an aGS interface 214 and a 10-GE interface 216, and a second pGS 220 having eight gigabit ports 222, an aGS interface 224 and a 10-GE interface 226. First pGS 210 is interconnected to second pGS 220 between 10-GE interface 214 and 10-GE interface 224. In this configuration, the interconnection of two pGSs results in a 16 port non-blocking solution.

[0057] FIG. 3 illustrates an fGS cross-bar fabric 300 having eight 10-GE interfaces 310. An fGS can be used to interconnect multiple devices together. The fGS cross-bar fabric illustrated in FIG. 3 can be used to interconnect eight pGS switches using the eight 10-GE interfaces 310.

[0058] FIG. 4 is an illustration of an fGS fabric 400 having eight 10-GE interfaces 410. Each of the eight pGS switches 420 each have eight gigabit ports 430, an aGS interface 440 and a 10-GE interface 450. Each of the eight pGS switches 420 are interconnected to one another through fGS fabric 400. Each of the 10-GE interfaces 450 of each of the pGS switches 420 are connected to an associated 10-GE interface 410 of fGS fabric 400. In this way the eight pGS switches 420 are interconnected through fGS fabric 400 by the 10-GE interfaces provided on each of the pGS switches 420 and the fGS fabric 400.

[0059] The aGS interfaces 440 are used to interconnect each of the pGS switches 420 to an external aGS for address resolution. However, each of the pGS switches 420 can in some embodiments of the invention perform address resolution functions without the use of an external aGS. For instance a pGS switch 420 could have an internal aGS. If pGS switch 420 does not detect an external aGS connected to aGS interface 440 the internal aGS will be used for address resolution. If pGS switch 420 does detect an external aGS connected to aGS interface 440 the external aGS will be utilized. An advantage of this feature is that it will allow for flexibility in the type of address resolution functions are performed. If the internal aGS is not desired an external aGS can be implemented providing more flexibility and modularity.

[0060] In one example of the present invention an FFP provided on pGS switch 420 has the functions of classification, metering, shaping and re-marking. In this example pGS switch 420 is provided with 16 filter masks and a 128K depth rules table. If an external aGS is connected to aGS interface 440, the internal FFP will not be utilized and the external aGS will be accessed. In this example the external aGS connected to aGS interface 440 provides 32 filter masks and a 256K rules table to pGS 420. Thus in this example, the same pGS 420 could be configured for 16 filter masks and a 128K rules table using the internal FFP or 32 filter masks and a 256K rules table using an external aGS through aGS interface 106 providing flexibility and modularity to pGS 420.

[0061] FIG. 5 is a flow diagram illustrating the flow of a packet received by a pGS. FIG. 5 is described in relation with FIG. 1. In step 510 a packet is received by pGS 100

through one of the gigabit ports 102. In one example of the invention, gigabit port 102 will have an ingress MAC for receiving the packet. In step 520 the packet will be processed by the ingress of the MAC and it will be determined if an external aGS is connected to aGS interface 106.

[0062] If an external aGS is detected, the packet is transmitted to aGS interface 106 in step 530. If an external aGS is not detected, the packet is processed using the internally provided aGS for address resolution. However, in some cases pGS 100 is not provided with an internal aGS for address resolution. In these cases step 520 is skipped and the packet is sent directly to aGS interface 106 as shown in step 530. When the packet is received by aGS interface 106 the packet is transmitted to the external aGS using a high speed serial interface or a parallel interface in order to transmit the packet externally as fast as possible so that minimal or no delay is experienced.

[0063] In step 540 the packet is received by the proper aGS and address resolution is performed on the packet. In some cases the address is learned and processed in order to transmit the packet to the proper port for output quickly and efficiently. In other cases the address is already learned and the packet can be transmitted directly to the proper output port. A variety of other address resolution functions can be performed by the aGS in order to help process and transmit the packet more efficiently and quickly. After the packet has been processed by the aGS, the packet is forwarded through pGS 100 for further processing. For example, in step 550 the packet can be sent to an MMU

for further processing and finally be transmitted to the proper output gigabit port 102 as shown in step 560. If the packet is being sent from an external aGS through aGS interface 106, the packet is sent to the MMU across a high speed serial interface or a parallel interface in order to transmit the packet externally as fast as possible so that minimal or no delay is experienced. If the packet is being transmitted internally a CPS bus can be used for transmission.

[0064] The above-discussed configuration of the invention is, in a preferred embodiment, embodied on a semiconductor substrate, such as silicon, with appropriate semiconductor manufacturing techniques and based upon a circuit layout which would, based upon the embodiments discussed above, be apparent to those skilled in the art. A person of skill in the art with respect to semiconductor design and manufacturing would be able to implement the various modules, interfaces, and tables, buffers, etc. of the present invention onto a single semiconductor substrate, based upon the architectural description discussed above. It would also be within the scope of the invention to implement the disclosed elements of the invention in discrete electronic components, thereby taking advantage of the functional aspects of the invention without maximizing the advantages through the use of a single semiconductor substrate.

[0065] Although the invention has been described based upon these preferred embodiments, it would be apparent to those of skilled in the art that certain modifications, variations, and alternative constructions would be apparent, while remaining within the spirit and scope of the invention. In order to determine the metes

and bounds of the invention, therefore, reference should be made to the appended claims.

FOUO " 6258860